

PRTS

Sec
Bib Data
insert
AI

~~Process for constructing a 3D scene model by analysing image sequences.~~

5 The invention relates to a process for constructing a 3D scene model by analysing image sequences.

The domain is that of the processing of image sequences and the modelling of real static scenes in a navigation context. The sequence consists of images relating to static scenes within which the viewpoint, that is to say the camera, changes.

10 The objective is to allow a user to navigate virtually in a real scene. However, the data regarding the scene consist of image sequences which may represent an enormous quantity of information. These sequences must be processed in such a way as to provide a compact representation of the scene, which can be used in an optimal manner for navigation, that is to say
15 allows interactive rendition, with controlled image quality. The problem is to obtain a high rate of compression whilst avoiding the techniques of inter-image predictive type which are not suited to navigation.

Description of Prior Art
Various representations of scenes currently exist. It is possible to distinguish mainly:

20 - representations based on 3D models, in which the geometry of the scene is generally represented in the form of plane facets with which texture images are associated. This modelling is much used to represent synthetic scenes obtained via software of the CAD (computer aided design) type. On the other hand, it is still little used to represent real scenes, since it is complex.
25 The current methods use few images, generally photographs, and the resulting representations are not very detailed and lack realism.

- non-3D representations obtained for example on the basis of the QuickTime VR software (Trademark of the Apple company). The data of the scene are acquired in the form of panoramic shots with transition image
30 sequences for switching from one panoramic shot to another. Such a representation considerably limits the possibilities of navigation in the virtual scene.

09831992-081601

0987199E 09760J

- 15

20

25

35

The invention also relates to a process of navigation in a 3D scene consisting in creating images as a function of the movement of the viewpoint,

characterized in that the images are created on the basis of the process for constructing the 3D model previously described.

The image sequences represent a very considerable quantity of data with high inter-image redundancy. The use of a 3D model which is the best model for representing a real static scene and the matching of the images via simple geometric transformations make it possible to broadly identify the inter-image redundancy. This model in fact makes it possible to take account of a large number of images. Moreover it requires no motion compensation operations at 2D image level.

A better compromise between compactness, that is to say compression of the data to be stored and processed, interactivity and quality of rendition is achieved: despite the high rate of compression, the process provides images of good quality and allows great flexibility and speed in navigation.

The invention makes it possible to obtain better realism than that obtained with the current 3D modelling techniques as well as better flexibility than that obtained with the conventional techniques for image coding.

Brief Description of the Drawings

The characteristics and advantages of the present invention will become more clearly apparent from the following description, given by way of example and with reference to the appended figures where:

- 25 - Figure 1 represents a processing algorithm describing the steps of
a process according to the invention,
 - Figure 2 represents the reference frames associated with a
viewpoint.

30 The acquisition of the data of the real scene is intimately related to the representation envisaged. In our example, we consider the situation where the images are acquired by a standard camera, at the video rate, and the camera movement is produced in a manner corresponding to the paths scheduled during utilization. In this context, the construction of a representation of a scene from image sequences may be likened to the techniques of image coding.

35 The principle of constructing the representation of a scene is to
select the necessary and sufficient data for reconstructing the images of the
sequence with controlled quality. The procedure consists in comparing the

09871992-081501a

images one by one so as to select the regions having the best relevance, a parameter which depends on the resolution and on the cost of description. In fact, the comparison is performed at the pixel level: the basic criterion for the comparison and selection of the pixels is the resolution of the 3D local surface associated with each pixel.

We assume that by suitable processing, known from the prior art, we obtain, for each viewpoint, its 3D position in a reference frame associated with the scene (position and orientation of the viewpoint), as well as a depth map associated with the image relating to the viewpoint. The object of the next phase is to construct a compact representation of all of these data which is suitable for navigation.

Figure 1 represents a flow chart describing the various steps of the process according to the invention.

At the system input, reference 1, we have data relating to an image sequence acquired by a camera moving within a real static scene as indicated earlier. However, it is entirely conceivable for certain moving objects to be present in the image. In this case, specific processing identifies these objects which are then marked so as to be ignored during subsequent processing. An ad hoc processing provides, for each image, a depth map as well as the position and the orientation of the corresponding viewpoint. There is no depth information in the zones corresponding to deleted moving objects.

A resolution value is calculated for each pixel of each image, this being step 2. A first and a second partitioning are then carried out during step 3. Step 4 performs a weight calculation for providing, step 5, relevance values allocated to the pixels. The next step 6 carries out a selection of the pixels depending their relevance. A sequence of masks of the selected pixels is then obtained for the image sequence, in step 7. After this step 7, steps 4 to 7 are repeated so as to refine the masks. These steps are repeated until the masks no longer change significantly. So then, step 8 is undertaken so as to carry out the construction of the faceted 3D model from the selected pixels alone.

Available at the system input, for each image of the sequence, is a
5 depth map as well as the position and the orientation of the corresponding
viewpoint.

The resolution at each pixel provides an indication of the level of detail of the surface such as it is viewed from the current viewpoint. It may be, for example, calculated over a block of points centred on the pixel and corresponds to the density of points in the scene, that is to say in 3D space, which relate to this block.

In one example, a window of 7x7 pixels, centred on the image pixel for which the resolution is calculated, is utilized. For each of the pixels belonging to this window, the depth information is processed so as to determine, from the distribution in 3D space of the points around the processed pixel, the 3D resolution: a distribution of the points over a large depth will give a less good resolution than a distribution of the points over a small depth. After processing all the pixels of the image, a resolution map of the image is obtained for each of the images of the sequence.

The navigation phase consists in interpolating the image of the current viewpoint from the 3D model. The model may be very large, and it is therefore useful to partition it so as to limit the quantity of information to be processed at each instant for the reconstruction of a viewpoint. Indeed, it is important for the images to be interpolated in a limited time so as to guarantee good fluidity of navigation. Moreover, the comparison of the images pixel by pixel in the data selection phase 6, described later, is an unwieldy operation, in particular if the sequences are long. This remark also holds for a partitioning, performed as early as possible, to reduce the quantity of calculations.

35

5 A first partitioning of the sequence is performed by identifying the viewpoints having no intersection of their observation fields. This will make it possible to avoid comparing them, that is to say comparing the images relating to these viewpoints, during subsequent steps. Any intersections between the observation fields, of pyramidal shape, of each viewpoint, are therefore determined by detecting the intersections between the edges of these fields. This operation does not depend on the content of the scene, but only on the relative position of the viewpoints. With each current image there is thus associated a set of images whose observation field possesses an intersection with that of this current image, this set constituting a list.

10 A projection is performed during this partitioning step 3 allowing a second partitioning. For each image group, a projection similar to that described later with regard to step 6, is carried out so as to identify the matching pixels. If an image has too few pixels matching with the pixels of an image of its list, this image is deleted from the list.

15 These partitionings, for each viewpoint, result in a list or group of viewpoints having 3D points in common with it, and which will therefore be compared during the selection of the pixels so as to reduce the redundancy. An array is constructed so as to identify, for each image of the sequence, the selected images required for its reconstruction.

20 During projection, the pixels having no match are marked by setting the resolution value, for example, to 1. By virtue of this particular marking, it will be evident, during step 6, that it is not necessary to re-project these pixels for the search for the matching pixels. This projection operation is in fact repeated in step 6 so as to avoid storing the information relating to these matches, obtained during step 3, this information representing a very large number of data.

25 Step 4 consists of a weight calculation for each of the pixels of an image. This parameter is introduced so as to take into account the cost of the pixels preserved. In the absence of any additional constraint on the selecting of the pixels, the latter may constitute regions of diverse sizes and diverse shapes and the cost of describing these regions may be high. To avoid this problem, a weight which takes into account the classification of the pixels in the close environment (pixel selected or not selected) is associated with each

35

09031902-001601

5

10

- 15

20

25

$$\text{relevance} = \text{resolution} \times (1 + \text{weight})$$

30

The selecting of the pixels is the subject of step 6

35

Here, for each pixel, this involves a search for the match in the other viewpoints, and involves a comparison of the relevance values for the identification of the pixel having best relevance.

5

10

15

20

25

30

35

relevance. However, while still reducing the redundancy, it may be advantageous to retain more of a representation of a given surface in order to avoid having to manipulate the representation at maximum resolution in order to create distant viewpoints. It is therefore advantageous to introduce a threshold into the comparison of the resolution values: if the ratio of two resolution values exceeds this threshold, none of the pixels is marked. Thus, each of the pixels can be used depending on the desired resolution, fine or coarse.

The marking of the pixels is done by firstly initializing all the pixels of all the masks, for example to the binary value 1. Each pixel is compared with its match, if it exists, in the other associated viewpoints during the partitioning phases. The one which possesses the lowest relevance is marked 0, that is to say it is rejected. Consequently, if none of its matches has a higher relevance than the current pixel, this is the one which is selected since it retains the initial marking. This therefore results, for each image of the sequence, in a binary mask or image, the pixels for which the value 1 is assigned corresponding to the selected pixels.

Step 7 collects the masks relating to each of the images forming the sequence in order to deliver the sequence of masks.

There is a feedback loop from step 7 to step 4 in order to refine the calculated relevance values. At each iteration, the weights and therefore the relevance values are recalculated from the masks obtained at the previous iteration.

The projection operations are repeated at each iteration and relate to all of the pixels of the image, pixels not selected during a previous iteration possibly being selected because, for example, of a reduction in the pertinence value of a pixel with which it is matched. However, the pixels not having a match in the other images are not projected.

To reduce the calculations, it is possible, at each iteration, to remove from the list of images which is associated with a current image the images no longer having a pixel with better relevance than the corresponding pixel in the current image. The final list of a given image thus contains the necessary and sufficient images for its reconstruction.

The iterative procedure is stopped after a predetermined number of iterations or when there are no longer any significant changes in the masks. Once these definitive masks have been obtained, step 8 follows step 7 and these masks are used in the phase of constructing the faceted 3D model, the

00001992-001601
100100-265F3860

The data relating to this faceted 3D model are composed of geometrical information and texture information. For each selected region, defined by the masks, its outline is polygonized and the corresponding depth map is approximated by 3D triangles. The selected texture data are grouped together so as not to retain unnecessary regions. A 3D model can easily be formed from all of this information. The list of the images and therefore the regions associated with each image can also be advantageously taken into account in the construction of the 3D model in order to partition it. This partitioning may then be used in the rendition phase in order to limit the amount of information to be processed during the image reconstruction.

The process of navigating in the 3D scene, which consists in creating images according to the movement of the viewpoint, uses all this information to recreate the images.